

Learning Execution Contexts from System Call Distributions for Intrusion Detection in Embedded Systems*

Man-Ki Yoon¹, Sibin Mohan², Jaesik Choi³, Mihai Christodorescu⁴ and Lui Sha¹

¹Department of Computer Science, University of Illinois at Urbana-Champaign

²Information Trust Institute, University of Illinois at Urbana-Champaign

³School of Electrical and Computer Engineering, Ulsan National Institute of Science and Technology

⁴Qualcomm Research Silicon Valley

Email: {mkyoon, sbin, lrs}@illinois.edu, jaesik@unist.ac.kr, mihai@qti.qualcomm.com

Abstract

Existing techniques used for intrusion detection do not fully utilize the intrinsic properties of embedded systems. In this paper, we propose a lightweight method for detecting anomalous executions using a *distribution of system call frequencies*. We use a cluster analysis to learn the legitimate execution contexts of embedded applications and then monitor them at run-time to capture abnormal executions. We also present an architectural framework with minor processor modifications to aid in this process. Our prototype shows that the proposed method can effectively detect anomalous executions without relying on sophisticated analyses or affecting the critical execution paths.

1 Introduction

An increasing number of attacks are targeting embedded systems [21, 37] that compromise the security, and hence safety, of such systems. It is not an easy task to retrofit embedded systems with security mechanisms that were developed for more general purpose scenarios since the former (a) have constraints in processing power, memory, battery life, *etc.* and (b) are required to meet stringent requirements such as timing constraints.

Traditional behavior-based intrusion detection systems (IDS) [10] rely on specific *signals* such as network traffic [15, 39], control flow [1, 8], system calls [14, 32, 5], *etc.* The use of system calls, especially in the form of sequences [14, 16, 44, 29, 12, 40], has been extensively studied in behavior-based IDSes for general purpose systems since many malicious activities often use system calls to execute privileged operations on system resources. Because server, desktop and mobile applications exhibit rich, wildly varying behaviors across executions, such IDSes need to rely either (a) on complex models of normal behavior, which are expensive to run and thus unsuitable for an embedded system, or (b) on simple, partial models, which validate only small windows of the application execution at a time. This opens the door for attacks where variations of a valid execution sequence are replayed with slightly different parameters to achieve a malicious goal; on the other hand, the application would not execute that sequence of operations in a normal manner, every time.

*This work is supported in part by grants from NSF CNS 13-02563, NSF CNS 12-19064, NSF CNS 14-23334 and Navy N00014-12-1-0046. Man-Ki Yoon was also supported by Qualcomm Innovation Fellowship and Intel PhD Fellowship. Jaesik Choi was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) grant funded by the Ministry of Science, ICT & Future Planning (MSIP) (No. NRF-2014R1A1A1002662 and No. NRF-2014M2A8A2074096). Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of sponsors.




Figure 1: An example system call frequency distribution (SCFD).

We observe that the very properties of embedded systems also make them amenable to the use of certain security mechanisms. The *regularity* in their execution patterns means that we can detect intrusions by monitoring the *behavior* of such applications [30, 45, 47, 46]; deviations from expected behavior can be considered to be malicious since the set of what constitutes legitimate behavior is often limited by design. In this paper we present an intrusion detection mechanism for embedded systems using a *system call frequency distribution* (SCFD). Figure 1 presents an example. It represents the *numbers of occurrences of each system call type* for each execution run of an application. The key idea is that the normal executions of an application whose behavior is predictable can be modeled by a small set of distinct system call distributions, each of which corresponds to a high-level *execution context*. We use a *cluster analysis* to learn *distinct* execution contexts from a set of SCFDs and to detect anomalous behavior using a *similarity metric* explained in Section 3.

Our detection method is lightweight, has a *deterministic* time complexity – hence, it fits well for resource-constrained embedded systems. This is due to the coarse-grained and concise representation of SCFDs. Although it can be implemented either at the operating system layer [33] or even used for offline analysis, we demonstrate an implementation on the SecureCore architecture [45] that increases security. In Section 4, we show that minor modifications to a modern multicore processor enables us to monitor and analyze the run-time system call usage of applications in a secure, non-intrusive manner.

We implemented our prototype on Simics, a full-system simulator [27]. Due to the inherent limitations in simulation environment, we developed a proof-of-concept implementation based on an example embedded application and various attack scenarios that highlight the mechanism and benefits of the SCFD-based intrusion detection and also its limitations. The experimental results show that SCFDs can effectively detect certain types of abnormal execution contexts that are difficult for traditional sequence-based approaches [6, 40]. Detailed results including a comparison with an existing sequence-based technique is presented in Sections 5 and 6.

Hence, the high level contributions of this paper are:

1. we introduce a lightweight method, utilizing the predictable nature of embedded system behaviors, with a deterministic time complexity for detecting anomalous execution contexts of embedded systems based on the *distribution of system call frequencies* (Section 3);
2. we present an architectural framework based on the SecureCore architecture for secure, non-intrusive monitoring and analysis of SCFDs (Section 4);
3. we demonstrate our techniques on a prototype implementation and evaluate its advantages and limitations using various attack scenarios that include a real attack (Sections 5 and 6).

2 Overview

The main idea behind SCFD is to learn the normal system call profiles, *i.e., patterns in system call frequency distributions*, collected during legitimate executions of a sanitized system. Analyzing profiles is

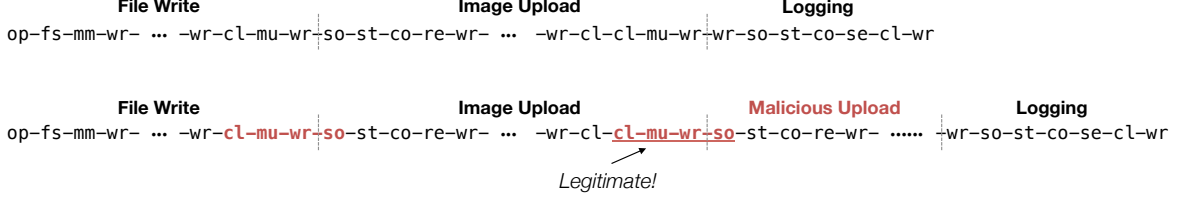


Figure 2: Sample sequence of system calls made by the target embedded application used in the evaluation (Section 5). The top sequence is from a normal execution. The bottom one is when a malicious code uses the exact same routine used by the normal execution for image uploading. The transition (that is underlined) is legitimate with respect to the set of legitimate sequences, and thus a sequence-based approach may fail to detect the malicious upload.²

challenging especially when such profiles change, often dramatically, depending on the execution modes and inputs. We address this issue by *clustering* the distribution of system calls capturing legitimate behavior. Each cluster then can be a *signature* that represents a high-level execution context, either in a specific mode or for similar input data. Then, given an observation at run-time, we test how similar it is to each previously calculated cluster. If there is no strong statistical evidence that it is a result of a specific execution context then we consider the execution to be malicious with respect to the learned model.

Attacks against sequence-based IDSes: Although sequence-based methods can capture detailed, temporal relations in system call usages, they may fail to detect abnormal *execution contexts*. This is because most sequence-based approaches rather profile the *local, temporal* relations among system calls within a limited time frame. Figure 2 highlights such a case. The sequence at the top is obtained from a part of normal execution of the target embedded application used in our evaluation (Section 5). A smart attacker may use the very same routine to circumvent the detection process and upload the image to its own server right after the normal image uploading operation completes, as shown at the bottom. A sequence-based method may not detect this malicious activity if the model parameter (*e.g.*, the sequence length or the Markovian order [6]) is not carefully chosen, since the transition `cl-mu-wr-so` is *not* abnormal with respect to what can be observed during normal executions.

In contrast to sequence-based techniques, our SCFD method may fail to detect a small local variation in system call sequences. However, as we show in this paper, it can easily detect abnormal deviations in high-level, naturally variable execution contexts such as the one illustrated above (Figure 2) since the SCFD significantly changes due to the malicious execution. Also, if the attacker corrupts the integrity of the data (for instance, replaces the input or changes its size to downgrade its quality) then our method is able to detect such problems – this is not easy for sequence-based methods as we explain in Section 6.3. Hence, by using these two approaches together, one can improve the overall accuracy of the system call-based IDS for embedded systems.

Assumptions and Adversary Model: The following assumptions are made in this paper:

1. We consider an embedded application that executes in a *repetitive fashion* – We monitor and perform a legitimacy test at the end of each invocation of a task.
2. We limit ourselves to applications of which most of the possible execution contexts can be profiled ahead of time. Hence, the behavior model is learned under the *stationarity* assumption – this is a general requirement of most behavior-based IDS. This can be justified by the fact that most

²Abbreviations: **a**ccess, **c**lose, **c**onnect, **e**xecve, **f**stat, **g**etuid, **m**map, **m**unmap, **o**pen, **r**ead, **s**endto, **s**ocket, **s**tat, **w**rite

embedded applications have a limited set of execution modes and input data fall within fairly narrow ranges. Also, a significant amount of analysis of embedded systems is carried out post-design/implementation anyways for a variety of reasons (to guarantee predictable behavior for instance). Hence, the information about the usage of system calls can be rolled into such a-priori analysis. Our method may not work well for applications that do not exhibit execution regularities (*e.g.*, due to frequent user interactions).

3. The initial state of the application is trustworthy. The profiling is carried out prior to system deployment. Also, any updates to the applications must be accompanied by a repeat of the profiling process. The application(s) could be compromised after the profiling stage, but we assume that the stored profile(s) cannot be tampered with. Again, such (repeat) analysis is typical in such systems – *e.g.*, anytime the system receives updates, including changes to the operating system or the processor architecture.
4. We consider threat models that involve changes to the behavior of system call usage. If an attack does not invoke or change any system calls, the activity at least has to affect executions afterward so that the future system call usage may change. The methods in this paper, as they stand, cannot detect attacks that never alter system call usage and that just replace certain system calls by *hijacking* them (*e.g.*, altering kernel system call table) [41].
5. We consider malicious code that can be secretly embedded in the application, either by remote attacks or during upgrades. The malicious code activates itself at some point *after* system initialization. We are not directly concerned with how the malicious code gained entry, but focus more on what happens after that.

As mentioned above, we assume that malware will exhibit a different pattern of system call usage. For example, malware that leaks out a sensitive information would make use of network-related system calls (*e.g.*, `socket`, `connect`, `write`, *etc.*) thus changing the frequencies of these calls.

3 Intrusion Detection Using Execution Contexts Learned from System Call Distributions

We now present our novel methods to detect abnormal execution contexts in embedded applications by monitoring changes in system call frequency distributions.

3.1 Definitions

Let $\mathcal{S} = \{s_1, s_2, \dots, s_D\}$ be the set of all system calls provided by an operating system, where s_d represents the system call of type d . During the n^{th} execution of an application, it calls a multiset σ^n of \mathcal{S} . Let us denote the n^{th} *system call frequency distribution* (or just *system call distribution*) as $\mathbf{x}^n = [m(\sigma^n, s_1), m(\sigma^n, s_2), \dots, m(\sigma^n, s_D)]^T$, where $m(\sigma^n, s_d)$ is the multiplicity of the system call of type d in σ^n . Hereafter, we simplify $m(\sigma^n, s_d)$ as x_d^n . Thus, $\mathbf{x}^n = [x_1^n, x_2^n, \dots, x_D^n]^T$.

We define a *training set*, *i.e.*, the execution profiles of a sanitized system, as a set of N system call frequency distributions collected from N executions, and is denoted by $\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^N]^T$. The clustering algorithm (Section 3.3) then maps each $\mathbf{x}^n \in \mathbb{N}^D$ to a cluster $c_i \in \mathcal{C} = \{c_1, c_2, \dots, c_k\}$. We denote by $c : \{\mathbf{x}^1, \dots, \mathbf{x}^N\} \rightarrow \mathcal{C}$ the cluster that $\mathbf{x}^n \in \mathbf{X}$ belongs to.

3.2 Learning a Single Execution Context

The variations in the usage of system calls will be limited if the application under monitoring has a simple execution context. In such a case, it is reasonable to consider that the executions follow a certain distribution of system call frequencies, clustered around a *centroid*, and cause a small variation from it

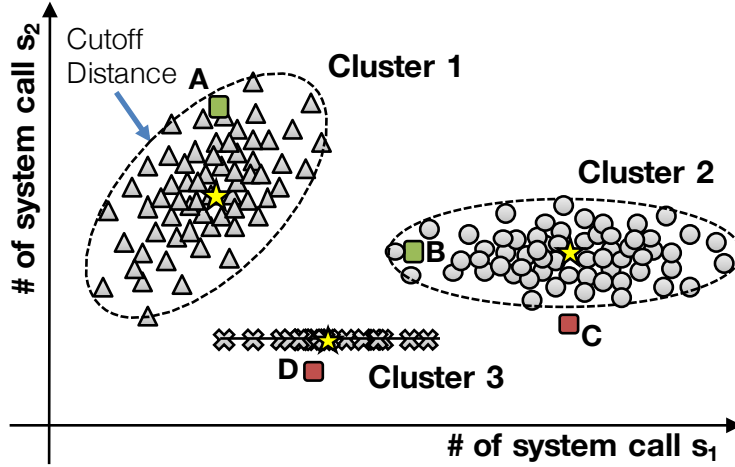


Figure 3: System call frequency distributions for $\mathcal{S} = \{s_1, s_2\}$ and clusters. The gray-colored objects are SCFDs in the training set. Each star-shaped point inside each cluster is its centroid. The ellipsoid around each cluster draws the cutoff line of the cluster; the points inside of the line are legitimate with respect to the cluster.

due to, for example, input data or execution flow. This is a valid model for many embedded systems since the code in such system tends to be fairly limited in what it can do. Hence, such analysis is quite powerful in detecting variations and thus catching intrusions.

For a multivariate distribution, the mean vector $\mu = [\mu_1, \mu_2, \dots, \mu_D]^T$, where $\mu_d = (\sum_n x_d^n)/N$, can be used as the centroid. Figure 3 plots the frequency distributions of two system call types (*i.e.*, $D = 2$). For now, let us consider only the data points (triangles) on the left-hand side of the graph. The data points are clustered around the star-shaped marker that indicates the centroid of the distribution formed by the points. Now, given a new observation from the monitoring phase, *e.g.*, the point marked ‘A’, a *legitimacy test* can be devised that tests *the likelihood* that such an observation is actually part of the expected execution context. This can be done by measuring *how far* the new observation is from the centroid. Here, the key consideration is on the *distance* measure for testing legitimacy.

One may use the Euclidean distance between the new observation \mathbf{x}^* and the mean vector of a cluster, *i.e.*, $\|\mathbf{x}^* - \mu\| = \sqrt{(\mathbf{x}^* - \mu)^T(\mathbf{x}^* - \mu)}$. Although the Euclidean distance (or L^2 -norm) is simple and straightforward to use, the distance is built on a strong assumption that each coordinate (dimension) contributes *equally* while computing the distance. In other words, the same amount of differences in x_1^n and x_2^n are considered equivalent even if, *e.g.*, a small variation in the usage of system call s_2 is the stronger indicator of abnormality than system call s_1 . Thus, it is more desirable to allow such a variable contribute more. For this reason, we use the *Mahalanobis distance* [28], defined as:

$$dist_M(\mathbf{x}^n, \mathbf{X}) = \sqrt{(\mathbf{x}^n - \mu)^T \Sigma^{-1} (\mathbf{x}^n - \mu)},$$

for a group of data set \mathbf{X} , where Σ is the covariance matrix of \mathbf{X} .³ Notice that the existence of Σ^{-1} is the necessary condition to define the Mahalanobis distance; *i.e.*, the difference of the frequency of each system call from the mean (*i.e.*, what is expected) is augmented by the *inverse of its variance*.

Accordingly, if we observe a small variance for certain system calls during the training, *e.g.*, `execve` or `socket`, we would expect to see a similar, small, variation in the usage of the system calls during

³ Σ is the positive definite. If we set $\Sigma = \mathbf{I}$, the Mahalanobis distance is equivalent to the Euclidean distance. Thus, the Mahalanobis distance is more expressive than the Euclidean distance.

actual executions as well. On the other hand, if the variance of a certain system call type is large, *e.g.*, **read** or **write**, the Mahalanobis distance metric gives a small weight to it in order to keep the distance (*i.e.*, abnormality) less sensitive to changes in such system calls. Cluster 2 in Figure 3 shows an example of the advantage of using the Mahalanobis distance over the Euclidean distance. Although C is closer to the centroid than B is in terms of the Euclidean distance, it is more reasonable to determine that C is an outlier and B is legitimate because we have not seen (during the normal executions) frequency distributions such as the one exhibited by C while we have seen a statistically meaningful amount of examples like B. As an extreme case, let us consider D which is quite close to Cluster 3’s center in terms of the Euclidean distance. However, it should be considered malicious because s_2 (*i.e.* the y -axis) should never vary.

Using covariance values also make it possible to learn *dependencies* among different system call types. For instance, an occurrence of the **socket** call usually accompanies **open** and many **read** or **write** calls. Thus, we can easily expect that changes in **socket**’s frequency would also lead to variations in the frequencies of **open**, **read** and **write**. Cluster 1 in Figure 3 is such an example that shows covariance between the two system call types. On the other hand, they are independent in Cluster 2 and 3. Thus, using the Mahalanobis distance we can not only learn how many occurrences of each individual system call should exist but also how they should vary together.

Now, given a set of system call distributions, $\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^N]^T$, we calculate the mean vector, μ , and the covariance matrix, Σ , for this data set. It then can be represented as a single cluster, c , whose centroid is defined as (μ, Σ) . Now, the Mahalanobis distance of a newly observed SCFD, \mathbf{x}^* , from the centroid is

$$dist(\mathbf{x}^*, c) = \sqrt{(\mathbf{x}^* - \mu)^T \Sigma^{-1} (\mathbf{x}^* - \mu)}. \quad (1)$$

If this distance is greater than a cutoff distance θ , we consider that the execution to be malicious. For example, B in Figure 3 is considered legitimate w.r.t. Cluster 2. One analytic way to derive this threshold, θ , is to think of the Mahalanobis distance w.r.t. the multinomial normal distribution,

$$p(\mathbf{x}^*) = \sqrt{|\Sigma|(2\pi)^D}^{-1} \exp\left(-\frac{1}{2}dist(\mathbf{x}^*, c)^2\right). \quad (2)$$

That is, we can choose a θ such that the p-value under the null hypothesis is less than a significant level p_0 , *e.g.*, 1% or 5%. Appendix A explains how to calculate θ given a p_0 .

3.3 Learning Multiple Execution Contexts Using Global k-means

In general, an application may show widely varying system call distributions due to multiple execution modes and varying inputs. In such scenarios, finding a single cluster/centroid for the whole set can result in inaccurate models because it would include many non-legitimate points that belong to none of the execution contexts – *i.e.*, the empty space between clusters in Figure 3. Thus, it is more desirable to consider that observations are generated from a set of *distinct* distributions, each of which corresponds to one or more execution contexts. Then, the legitimacy test for a new observation \mathbf{x}^* is reduced to identifying the *most probable* cluster that may have generated \mathbf{x}^* . If there is no strong evidence that \mathbf{x}^* is a result of an execution corresponding to any cluster then we determine that \mathbf{x}^* is most likely due to malicious execution.

Suppose we collect a training set $\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^N]^T$ where $\mathbf{x}^n \in \mathbb{N}^D$. To learn the distinct distributions, we use the k -means algorithm [24] to partition the N data points on a D -dimensional space into k clusters. The k -means algorithm works as follows:

1. Initialization: Create k initial clusters by picking k random data points from \mathbf{X} .
2. Assignment: For each $\mathbf{x}^n \in \mathbf{X}$, assign it to the closest cluster $c(\mathbf{x}^n)$, *i.e.*,

$$c(\mathbf{x}^n) = \arg \min_{c_k \in \mathcal{C}} dist(\mathbf{x}^n, c_k). \quad (3)$$

ALGORITHM 1: GLOBAL K-MEANS(\mathbf{X} , MAX_K , Bound_{TD})

```
1:  $\{\mathbf{X}$ : the training set $\}$ 
2:  $\{\text{MAX}_K$ : the maximum number of clusters $\}$ 
3:  $\{\text{Bound}_{TD}$ : the total distance bound $\}$ 
4: Create  $c_1$  with  $\mathbf{X}$ . Calculate  $\mu_1$  and  $\Sigma_1$ .
5:  $\mathcal{C} \leftarrow \{c_1\}$ 
6:  $k \leftarrow 2$ 
7:  $\text{Min}_{TD} \leftarrow \infty$ 
8: while  $k \leq \text{MAX}_K$  or  $\text{total-dist}(\mathbf{X}, \mathcal{C}) > \text{Bound}_{TD}$  do
9:   for  $n = 1, \dots, N$  do
10:    Create  $c_k$  with  $\mathbf{x}^n$  as its initial point.
11:     $\mathcal{C}' \leftarrow \text{k-means}(\mathbf{X}, \mathcal{C} \cup c_k)$ 
12:    if  $\text{total-dist}(\mathbf{X}, \mathcal{C}') < \text{Min}_{TD}$  then
13:      {Note: The best clustering for  $k$  so far}
14:       $\mathcal{C}^* \leftarrow \mathcal{C}'$ 
15:       $\text{Min}_{TD} \leftarrow \text{total-dist}(\mathbf{X}, \mathcal{C}')$ 
16:    end if
17:  end for
18:   $\mathcal{C} \leftarrow \mathcal{C}^*$ 
19:   $k \leftarrow k + 1$ 
20: end while
21: return  $\mathcal{C}$ 
```

3. Update: Re-compute the centroid (*i.e.*, μ and Σ) of each cluster based on the new assignments.

The algorithm repeats steps (2) and (3) until the assignments stop changing. Intuitively speaking, the algorithm keeps updating the k centroids until the total distance of each point \mathbf{x}^n to its cluster,

$$\text{total-dist}(\mathbf{X}, \mathcal{C}) = \sum_{n=1}^N \text{dist}(\mathbf{x}^n, c(\mathbf{x}^n)), \quad (4)$$

is minimized.

The k -means algorithm requires a strong assumption that we already know k , the number of clusters. However, this assumption does not hold in reality because the number of distinct execution contexts is not known ahead of time. Moreover, the accuracy of the final model heavily depends on the initial clusters chosen randomly.⁴ Hence, we use the *global k -means* method [23] to find the number of clusters as well as the initial assignments that lead to *deterministic accuracy*. Algorithm 1 illustrates the global k -means algorithm. Given a training set \mathbf{X} of N system call frequency distributions, the algorithm finds the best number of clusters and assignments. This is an incremental learning algorithm that starts from a single cluster, c_1 , consisting of the entire data set. In the case of $k = 2$, the algorithm considers each $\mathbf{x}^n \in \mathbf{X}$ as the initial point for c_2 and runs the assignment and updates steps of k -means algorithm. After N trials, we select the final centroids that resulted in the smallest total distance calculated by Eq. (4). These two centroids are then used as the initial points for the two clusters, respectively, in the case of $k = 3$. This procedure repeats until either k reaches a pre-defined MAX_K , the maximum number of clusters, or the total distance value becomes less than the total distance bound Bound_{TD} . Note that the total distance in Eq. (4) decreases monotonically with the number of clusters. For example, if every point is its own cluster then the total distance is zero since each point itself is the centroid.

⁴Finding the optimal assignment in the k -means algorithm with the Euclidean distance is NP-hard. Thus, finding the optimal assignments with the Mahalanobis distance is at least NP-hard because the Mahalanobis distance is more general than the Euclidean distance.

The original algorithm assumes the Euclidean distance. As explained above, we use the Mahalanobis distance as in Eq. (1). Meanwhile, $k\text{-means}(\mathbf{X}, \mathcal{C})$ (line 11) is the standard k -means algorithm without the random initialization; it assigns the points in \mathbf{X} to a $c_k \in \mathcal{C}$, update the centroids, repeats until stops, and then returns the clusters with the *updated* centroids. The standard k -means algorithm uses the Euclidean distance and thus the centroids of the initial clusters are the data points that were picked first. Remember, however, that the Mahalanobis distance requires a covariance matrix. Since there would be only one data point in each initial cluster we use the global covariance matrix of the entire data set \mathbf{X} for the initial clusters. After the first iteration, however, the covariance matrix of each cluster is updated using the data points assigned to it.

The clustering algorithm finally assigns each data point in the training set into a cluster. Then, each cluster $c_i \in \mathcal{C}$ can be represented by the centroid, (μ_i, Σ_i) , that now makes it possible to calculate the Mahalanobis distance of a newly observed SCFD \mathbf{x}^* to each cluster using Eq. (1). The legitimacy test of \mathbf{x}^* is then performed by finding the closest cluster, c^* , using Eq. (3). Thus, if

$$\text{dist}(\mathbf{x}^*, c^*) = \min_{c_i \in \mathcal{C}} \text{dist}(\mathbf{x}^*, c_i) > \theta$$

for a given threshold θ , we determine that the execution does not fall into any of the execution contexts specified by the clusters since $\text{dist}(\mathbf{x}^*, c_i) > \theta$ for all $i = 1, \dots, k$. We then consider the execution to be malicious. As an example, for the new observation C in Figure 3, Cluster 2 is the closest one and C is outside its cutoff distance. Thus, we consider that C is malicious. Note that, as shown in the figure, the same cutoff distance defines different ellipsoids for different clusters; each ellipsoid is a equidistant line from the mean vector measured in terms of the *Mahalanobis* distance. Thus, a cluster with small variances (*i.e.*, less varying execution context) would have a smaller ellipsoid in the Euclidean space.

3.4 Dimensionality Reduction

The number of system call types, *i.e.*, D , is quite large in general. Thus, the matrix calculations in Eq. (1) might result in an unacceptable amount of analysis overhead.⁵ However, embedded applications normally use a limited subset of system calls. Furthermore, we can significantly reduce the dimensionality by ignoring system call types that never vary. Consider Cluster 3 from Figure 3. Here, x_2 can be ignored since we can reasonably expect it to never vary during the normal execution.⁶ Thus, before running the clustering algorithm, we reduce \mathcal{S} to $\mathcal{S}' = \{s_{d_1}, s_{d_2}, \dots, s_{D'}\}$, where $D' \leq D$, such that the variance of x_d for each $s_d \in \mathcal{S}'$ is non-zero in the entire training set \mathbf{X} . However, we should still be able to detect any changes in such system calls that never varied (including those that never appeared). Thus, we merge all such x_d in $\mathcal{S} - \mathcal{S}'$; the sum should not change in normal executions. In case D' is still large, one may apply a statistical dimensionality reduction technique such as Principal Component Analysis (PCA) [18] or Linear Discriminant Analysis (LDA) [13].

4 Architectural Support for SCFD Monitoring

Most the existing system call-based intrusion detection systems rely on the operating system to provide the information, say, by use of auditing modules [38, 33]. While this provides the ability to monitor extensive resources such as system call arguments, it requires the operating system itself to be trustworthy. In this paper we avoid this problem by proposing a new architectural framework that requires minor micro-architecture modifications. The architecture builds upon the SecureCore architecture [45] that enables a trusted on-chip entity, *e.g.*, a *secure core*, to continuously monitor the run-time behavior of applications on another, potentially untrusted entity, the *monitored core*, in a *non-intrusive manner*.

⁵Note that μ and Σ values are calculated from the clustering algorithm which is an *offline* analysis. Thus we store Σ^{-1} for computational efficiency.

⁶In fact, s_2 cannot be ignored in the example depicted in Figure 3 since its variance is non-zero in clusters 1 and 2.




Figure 4: The SecureCore architecture for SCFD monitoring. The SCTM traces the system calls made by applications running on the monitored core. The traces, written in the SPM, are retrieved by the secure monitor to perform legitimacy test for the executions.

In this section we describe the modifications to SecureCore that enable us to monitor system calls. We refer interested readers to [45] for the full details about the SecureCore architecture.

4.1 Overview

Figure 4 shows the overall architecture for system call monitoring. It consists of (a) a secure core, (b) a monitored core, (c) an on-chip system call tracing module (SCTM) and (d) a scratch pad memory (SPM). The secure core uses the SCTM and SPM to monitor the usage of system calls by applications executing on the monitored core. The SCTM extracts relevant information from the monitored core and then writes it to the SPM. A monitoring process on the secure core then uses this information to check whether the run-time behavior has deviated from the expected behavior that we profiled using the method described in Section 3.

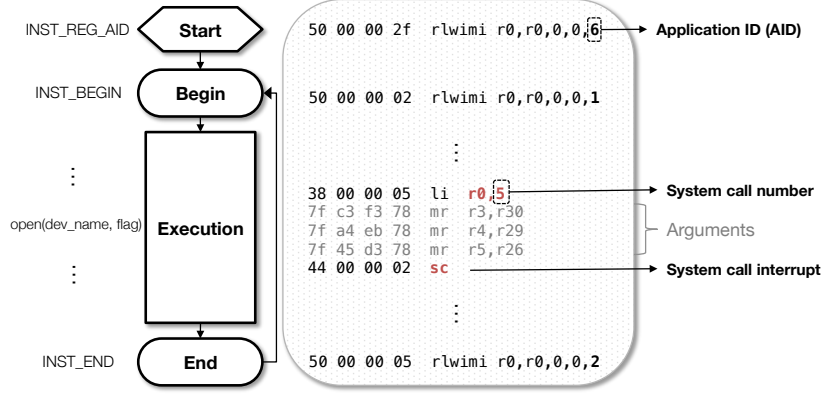
Note: We capture the profile of *normal* executions in a similar manner: the monitoring process collects SCFDs using the SCTM and SPM under trusted conditions. We then apply the learning algorithm from Section 3. The resulting normal profile (one per application) is then stored in a secure memory location.

4.2 System Call Tracing Module (SCTM)


The system call tracing module (SCTM) tracks how many times each application on the monitored core uses each system call type (*i.e.*, SCFDs). The main point is to catch the moment each call is invoked. We are able to do this because, in most processor architectures, a specific instruction is designated for this very purpose, *i.e.*, for triggering system calls. The calling conventions vary across processor architectures and operating systems. In the PowerPC architecture, that our prototype is based on, an `sc` instruction issues a system call based on a number stored in the `r0` register [11]. The actual call is then handled by the operating system kernel. Hence, the execution of the `sc` instruction denotes the invocation of a system call.

Another piece of information that we require is *who* initiated the call; hence we introduce a new instruction to help with identifying the requester. Figure 5(a) describes the process by which the SCTM gathers the required information from an application. When an application starts, it registers its *Application ID* (AID) and *Process ID* with the SCTM. Here, the AID is a unique numerical value assigned to each application. It is used to let the secure monitor (on the secure core) locate the correct profile. Once the registration is complete, the secure monitor is able to map each application to corresponding PIDs. The above registration process is carried out by a special instruction, `INST_REG_AID`, as described in the figure.⁷ The special instruction has other modes as well: (i) `INST_BEGIN` and (ii) `INST_END` that

⁷`rlwimi` instruction is the *Rotate Left Word Immediate Then Mask Insert* instruction in the PowerPC ISA [11]. An execution of `rlwimi 0,0,0,0,i` for $0 \leq i \leq 31$ is equivalent to a `nop` – hence, we used it for our purposes.



(a) System call catch by SCTM



(b) SCFD update in the scratch pad memory (SPM)

Figure 5: (a) The System Call Tracing Module (SCTM) catches system call executions by looking at the instruction decoder. (b) Upon execution of a system call, the corresponding counter (e.g., Cnt5 for `open`) of the SCFD entry mapped to the application is incremented. The gray areas are for alignment padding.

demarcate the region where we are tracking the usage of system calls. Once `INST_END` completes, the monitor retrieves the data collected from the recently completed region of execution and applies the detection algorithm. The data is reset with the execution of an `INST_BEGIN`. While an attacker may try to execute malicious code block before `BEGIN` or after `END` to avoid detection, we can catch such situations because there should be no system call execution during that point in the code. Thus, in such cases, the SCTM would immediately raise an alarm. Also, an attack may skip some of all of the special instructions or modify any of the values. Again, these cannot help the attacker hide malicious code execution because the system call distribution would need to be consistent with the profile. Also, watchdog timers can be used to check whether the applications are executing the special instructions in time.

Figure 5(a) shows how an `open` system call invocation is detected. As explained above, the system call number, 5 in Linux, is written to the `r0` register. Thus, by looking at the value of the register, we can track which system call is being invoked. When an `sc` instruction is executed, SCTM takes the PID register and the `r0` register values. It then updates the corresponding *SCFD entry* in the scratch pad memory (see Figure 5(b)). An entry is a contiguous memory region of length $2D + 4$ bytes, where D is the total number of system calls provided by the OS. Using the PID, SCTM locates the corresponding entry and then increments the *counter* of system call d if the value in `r0` was d . The sizes of the SPM and each entry field are implementation-dependent. In our implementation, we assume at most 382 system call types (that are enough to cover most Linux implementations) that results in the size of an entry being 768 bytes at most. Thus, an SPM of size 8KB can provide a space for around 10 applications to be monitored simultaneously. The SPM can be accessed only by the secure core. When an `INST_END`




Figure 6: The system setup for evaluation. The modified processor architecture (Figure 4 in Section 4) is implemented on Simics.

is executed, the secure monitor reads the corresponding entry from the SPM,⁸ a new SCFD, finds the profile for the corresponding application using the PID to AID map, then verifies the legitimacy of the SCFD.

5 Evaluation Framework

In this section, we first present the implementation details for our prototype (Section 5.1), the application model for our experiments (Section 5.2) and some attack scenarios that are relevant to this application (Section 5.3).

5.1 System Implementation

We implemented a prototype of our SCFD-based intrusion detection system on Simics [27]. Simics is a full-system simulator that can emulate a hardware platform including real firmware and device drivers, and also allows for processor micro-architecture modifications. Figure 6 shows the system setup used for our evaluation. We used the Freescale MPC8641D [31] development platform on Simics. It has a dual-core processor, each core of which runs at 1350MHz and the system has a memory of 1GB and runs an embedded Linux 2.6.23. The SCTM was implemented by extending the `sample-user-decoder` on Simics. This allows us to implement the necessary ISA modification as described in Section 4.2. The SPM has a total size of 8KB.

5.2 Target Application Model

Figure 7(a) shows the target application. Each invocation of the application (period of one second) cycles through the following steps: (i) retrieve a raw image from a camera, (ii) compresses it to a JPEG format, (iii) upload the image file to the base station through FTP and finally (iv) write a log via HTTP post. This type of application model (image capture → processing → communication) can be found in modern unmanned aerial vehicles (UAVs) that are used for surveillance or environmental studies [19].

⁸An interrupt can be raised by SCTM to inform the secure monitor of the execution of `INST_END`. However, if possible, it should be avoided because the secure core can be continuously interrupted by a compromised application on the monitored core thus degrading its ability to perform the monitoring. If necessary, the SCTM should be configured to block consecutive `INST_END` instructions. Of course, it is more preferable for the monitor to poll the SCTM.

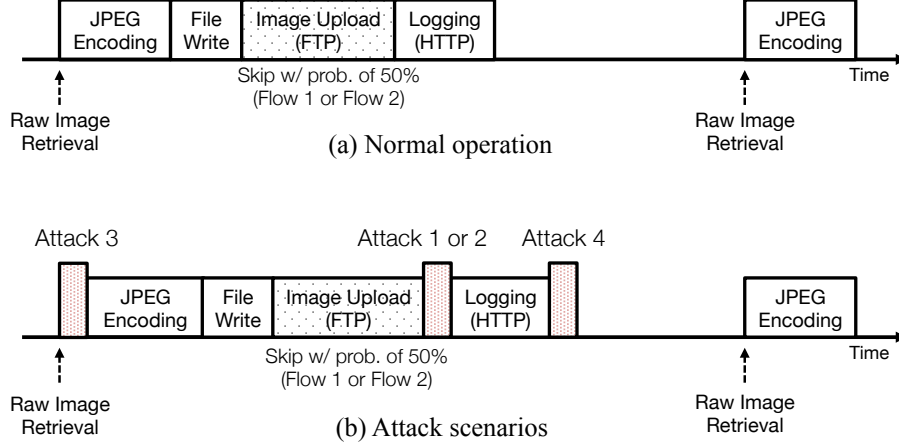


Figure 7: The normal execution flow of the target application. **Attack 1** and **2** leak user authentication information and JPEG image through HTTP and FTP, resp., **Attack 3** erases the raw image buffer and **Attack 4** is a real-world shellcode that spawns a shell (`/bin/sh`) by executing `execve` [36].

The distributions of the system call frequencies exhibited by this application is mainly affected by the stages after the JPEG compression. While the raw image size is always fixed (*e.g.*, $2.6MB$ for 1280×720 resolution), a JPEG image size can vary ($27KB - 97KB$) because of compression. This results in a variance in the number of `read` and `write` system calls. To increase the complexity of the application and also to show certain scenarios that the proposed method cannot deal with well, we added an additional code branch before the FTP upload stage that behaves as follows: the system can randomly skip the image upload process (based on a probability of 0.5). This affects the number of occurrences of network and file-related system calls during actual execution. Hence, the application has two legitimate flows, **Flow 1** and **Flow 2** as shown in Figure 8; the figure also shows the system call types used at each stage of the execution flows.

We use this type of application model for the following reasons: (a) Simics, being a full system simulator that executes on a ‘host’ system, is not fast enough to be able to control an actual system. Hence, we need to develop an application model that it can simulate; (b) we still need to demonstrate, in the simplest possible way, how our SCFD-based intrusion detection system works – this application model is able to highlight the exact mechanisms and even its limitations. Note that this target application is crafted to show *more* variance than many real embedded systems. Hence, if our detection method can catch changes in the system call distributions here then it can detect similar attacks in embedded systems that show less variance.

5.3 Attack Scenarios

We consider the following attack scenarios for this application:

1. **Attack 1** steals user authentication information used to connect to the base station’s FTP server and sends it to an adversary HTTP server. *This attack invokes the same HTTP logging calls used by the legitimate executions.*
2. **Attack 2** uploads the image that was just encoded by the application to an adversary FTP server. *This attack also uses the same functions used by a legitimate FTP upload.*
3. **Attack 3** modifies the raw image array received from the camera. The attack erases the array by calling `memset`. *This attack does not require any system calls.*

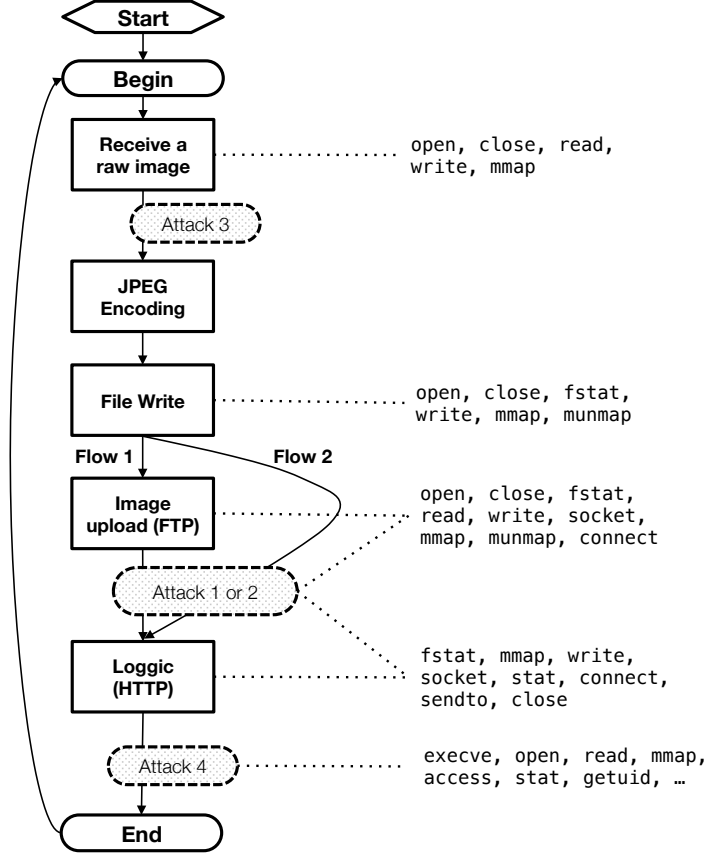


Figure 8: The execution flow of the target application and the system call types used at each stage. Apart from those shown in the figure, the application used `futex`, `rt_sigreturn` and `brk` system calls.

4. **Attack 4** is a real shellcode targeted at Linux on PowerPC and executes `execve` to spawn a shell (`/bin/sh`) [36]. In general, a shellcode can be injected by data sent over a network or from a file and can be executed by exploiting buffer overflow or format string vulnerabilities. In our implementation, the shellcode is stored in `char shellcode[]` and is executed by `__asm__("b shellcode")` when enabled.

The attack codes execute at spots marked in Figure 7 when enabled. Note that our method is independent of where they happen since SCFDs do not care about the sequences of system calls.

6 Evaluation Results

We now evaluate the SCFD method on the prototype described in the previous section. We also compare it with an existing sequence-based approach to show how the two methods can be used to complement each other.

6.1 Training

To obtain the training set, we executed the system under normal conditions (*i.e.*, no attack present) 2,000 times. The target application used 14 types of system calls (as shown in Figure 8, together with `futex`, `rt_sigreturn` and `brk`). We used the learning algorithm presented in Section 3 with settings,

Table 1: The mean and the standard deviation of the system call frequency distributions in the entire training set and in each cluster after running the learning method. The boldfaced values represent the system call types whose variance is non-zero.

	# pts		write	read	mmap	open	close	fstat	munmap	socket	connect	stat	Execution context
All	2000	Mean	29.519	101.197	1.520	2.514	4.548	1.520	1.520	2.034	2.034	4.034	
		Stdev	10.602	10.135	0.500	0.500	1.496	0.500	0.500	0.997	0.997	0.998	
c_1	490	Mean	17.376	91.000	1.000	2.000	3.000	1.000	1.000	1.000	1.000	3.000	Small image
		Stdev	1.246	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	No FTP upload
c_2	519	Mean	33.613	108.306	2.000	3.000	6.000	2.000	2.000	3.000	3.000	5.000	Small-medium image
		Stdev	2.539	1.269	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	FTP upload
c_3	506	Mean	43.708	113.354	2.000	3.000	6.000	2.000	2.000	3.000	3.000	5.000	Medium-large image
		Stdev	4.539	2.269	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	FTP upload
c_4	335	Mean	21.176	91.000	1.000	2.000	3.000	1.000	1.000	1.000	1.000	2.998	Medium image
		Stdev	1.080	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.055	No FTP upload
c_5	150	Mean	25.575	91.000	1.000	2.000	3.000	1.000	1.000	1.000	1.000	3.000	Large image
		Stdev	1.627	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	No FTP upload

$\text{MAX}_k = 10$ and the total distance bound Bound_{TD} of 1,000 on the resulting traces. The cutoff distance, θ , which attests to the legitimacy of a new SCFD, is 1.95996; *i.e.*, the significance level is 5%. We also tested for $p_0 = 1\%$, *i.e.*, $\theta = 2.57583$. Again, the lower significance level, the more we are confident about the statistical significance when an outlier is observed.

Table 1 summarizes the training results. The first row shows the mean and the standard deviation of the SCFDs in the whole training set. The algorithm first reduce the dimensionality of the results from 14 to 10 by removing the system call types that show zero variance in the training set. The variations of **write** and **read** are due to the JPEG compression and the FTP uploading phases. The latter also affects the frequencies of the network- and file-related system call types. The global k -means algorithm stopped at $k = 5$ (the moment when the total distance becomes less than the bound Bound_{TD}) resulting in *five clusters* as shown in the same table.

From these results, note first that the variation of each system call type is significantly reduced after the clustering; most of them become zero. This is because each cluster contains similar SCFDs, representing similar execution contexts. Also, from observing the mean values of the system call types other than **write** and **read**, we can infer that Clusters c_2 and c_3 are from a similar execution context while the others are from a different context. We also observe that the former group corresponds to Flow 1 because of the additional system calls required for the FTP transfer. Also, the fewer number of **write** and **read** system calls of the second group suggest that they belong to Flow 2. As expected, within each group clusters are distinguished by **write** and **read** due to the varying sizes of images that are compressed. The clustering results would be similar if MAX_k was set to, for example, 2. In this case, one cluster would have the points from Clusters c_1 , c_4 and c_5 combined but with different centroid and similarly c_2 and c_3 would constitute another new cluster. This could, however, blur boundaries between the execution contexts.

6.2 Accuracy

Now, we evaluate the *accuracy* of our intrusion detection methods. We enabled each of the attacks from Section 5.3. For each attack type, we carried out 300 execution instances and measured how many times the monitor detects malicious execution. An execution is considered malicious if any of the following is true: (i) any system call other than the 14 observed types is detected; (ii) any system call whose variance was zero during the profile (4 out of 14 in the case above) actually exhibits variance or (iii) the distance of an observation from its closest cluster is longer than the threshold. Among these, rule (i) was never observed in the cases of Attacks 1-3 because Attacks 1 & 2 re-used the same functions from normal executions and Attack 3 makes no system calls at all. Table 2 summaries the results of our

detection method as well as those of the sequence-based approach explained in Section 6.3. The results of our SCFD method are as follows.

1. **Attack 1 (HTTP post):** *All* the executions were classified as malicious based on rule (ii) above since one additional `brk` and `sendto`, each, were invoked. We tested the executions again, after removing such obvious situations. The results, however, did not change – *all the malicious executions were caught by our monitor*. Because of the additional HTTP request by the attack code, `socket`, `connect`, `close` and `stat` were called more in both Flows 1 and 2. System call types other than the ones mentioned here were consistent with the profile. With $p_0 = 1\%$, the results were the same since the use of the additional system calls already increased the distances from clusters to lie outside the acceptable boundaries.
2. **Attack 2 (FTP upload):** If the attack code executes on Flow 1, it is easily caught because of the additional FTP transfer. As with the above, this attack changes some network related system calls. This is enough to make the SCFDs fall outside the legitimate regions. The attack reads the image file as well. This increases the usage of `read` system calls thus further highlighting the anomalous behavior.

On the other hand, if the attack is launched on Flow 2 (that skips the FTP upload to the base station), it may not be as easy to detect. Since the attack uses the same functions that are invoked by legitimate code, it looks like the application is following Flow 1 (where the FTP upload is actually legitimate). In this case, only 1% of the malicious code executions were caught. The detection rates would be significantly higher if the attacker either used different images that vary in size or used code that utilizes different combinations of system calls. The latter case would also hold for the HTTP post attack. **Note:** the detection was not successful because we *tailored* the attack instance to *closely match* legitimate execution (especially due to our knowledge of the detection methods); however, many attacks will not be able to match legitimate execution in such a precise manner and will end up being caught.

3. **Attack 3 (Data corruption):** This attack does not use any system calls; it just changes the values of the data. However, this may affect the execution of code segments that follow, especially ones that depend on the data – the JPEG compression. The attack code resets the raw image data by using `memset`. This is compressed by the JPEG encoder that produces 15 KB of black images. This attack was *always* caught by our monitor because these image sizes are not typical during normal execution. Hence, calls to `read` and `write` were much less frequent when compared to normal execution. The attack could have circumvented our detection method if, *e.g.*, the raw image is just replaced with another that has a similar after-compression size as the original or only a part of the image is modified.⁹ But performing either of these actions may also trigger the use of additional system calls that would be caught by our monitor.
4. **Attack 4 (Shellcode execution):** This attack was easily detected since it uses `execve` (used by the shellcode) which was never observed during the profiling phase. Furthermore, it was followed by a bunch of system calls including `open`, `mmap`, `access`, `getuid`, *etc.* This was due to the execution of a shell, `/bin/sh`, spawned by the injected `execve`. In fact, `INST_END` was not executed since `execve` does not return on success. Nevertheless, the attack could be detected because a watchdog timer was used to *wake up* the secure monitor that then checks the application’s SCFD traced until the timer expires. From this experiment it can be expected that our method can detect more sophisticated shellcode [35] that uses unusual system calls, *e.g.*, `setreuid`, `setregid`, *etc.*¹⁰

⁹If the attacker modified the compressed image, our method cannot detect it because the system call usage would never change. This, however, does not fall into our threat model explained in Section 1.

¹⁰The shellcode used in our experiment is simpler than the ones targeted for Linux/x86 [35] due to the scarcity of sophisticated shellcode for Linux/PowerPC.

Table 2: Comparison between SCFD and PST (a sequence-based) methods.

Type	SCFD	PST	Cause
Attack 1	100%	100%	SCFD: Extra network-related system calls PST: Unusual transition (HTTP-HTTP)
Attack 2 (Flow 1)	100%	0% ($\bar{N} = 3$) 100% ($\bar{N} = 5$)	SCFD: Extra network- and file-related system calls PST : Short sequence cannot capture unusual transition (FTP-FTP)
Attack 2 (Flow 2)	1%	0% (both \bar{N})	Both: Not differentiable from legitimate Flow 1
Attack 3	100%	0% ($\bar{N} = 3$) 100% ($\bar{N} = 5$)	SCFD: Too small image size PST: Short sequence cannot capture shortened <code>write</code> chain
Attack 4	100%	100%	Both: <code>execve</code> was never seen

False Positives: The *false positive rate* is just as important as the detection rate because frequent false alarms degrade system availability. To measure the false positive rates, we obtained a new set of SCFDs by running the system without activating any attacks and measured how many times the secure monitor classifies an execution as being suspicious. Most false positives in these tests were due to the images sizes that, when compressed, fell below the normal ranges. For the cut-off distance θ with $p_0 = 5\%$, 35 out of 2,000 executions (1.75%) were classified as malicious. With $p_0 = 1\%$, *i.e.*, a farther cut-off distance, it was reduced to just 17 (0.85%). Such a lower significant level relaxes the cut-off distance and produces fewer false alarms because even some rarely-seen data points are considered normal. However, this may result in lower detection rates as well. In the attack scenarios listed above, however, the results did not change even with $p_0 = 1\%$. This is a consideration for system designers to take into account when implementing our intrusion detection methods; they will have a better feel for when certain executions are normal and when some are not. Hence, they can decide to adjust values for p_0 based on the actual system(s) being monitored.

These results show that our method can effectively detect malicious execution contexts without relying on complex analysis. While it is true that the accuracy of the method may depend on the attacks that are launched against the system, in reality an attacker would need to not only know the exact distributions of system call frequencies but also be able to implement an attack with such a limited set of calls – both of these requirements significantly raise the difficulty levels for would-be attackers.

6.3 Comparison with Sequence-based Approach

To show how our detection method can effectively complement existing system call-based intrusion detection methods, we compare it against a sequence-based approach using the same data set used in the previous section. Among the existing methods (explained in Section 7), we use a variable-order Markovian model (VMM) [2] using *Probabilistic Suffix Tree* (PST) [40, 6], due to its ability to learn significant sequence patterns of different lengths. This enables us to calculate $\Pr(s_{(t)} | s_{(t-N)} \cdots s_{(t-2)} s_{(t-1)})$, that is, the probability of a system call made at time t given a recent history, without having to learn all or a fixed-length of the sequences (used by N -gram or fixed-order Markovian models). The sequence length, N , varies with different patterns and is learned from their significance in the training set.

PST learns the conditional probabilities in a (suffix) tree structure and thus requires a user defined parameter \bar{N} that limits the maximum depth of the tree, *i.e.*, the length of sequence pattern. We tested two configurations: (i) $\bar{N} = 3$ and (ii) $\bar{N} = 5$ that show different results. Given the PST learned from the training set, a test is carried out as follows: for each system call, we calculate its (conditional)

probability using the PST and consider it malicious if the probability is less than a threshold of 1%. An iteration (*i.e.*, one execution) of the target application is classified as being malicious if any call is classified as malicious. Table 2 summaries the results from the detection techniques.

1. **Attack 1 (HTTP post):** PST was able to detect Attack 1 because of this particular sequence: `sendto-close-write-write`, where we have that $\Pr(\text{wr}|\text{se} - \text{cl} - \text{wr}) = 0$ since in the legitimate executions, `sendto-close-write`, the end of the HTTP logging function, is almost always followed by `stat`. The last `write`, which is the beginning of the HTTP function caused an *unusual transition* and hence detected by PST. On the other hand, SCFD detected the attack because of the *unusual frequencies* of network-related system calls.
2. **Attack 2 (FTP upload):**¹¹ In normal executions of Flow 1, the legitimate FTP operation ends with this sequence: `write-close-close-munmap-write`. Then, the legitimate HTTP operation starts with a `write` call. The extra FTP operation by Attack 2, which executes between the two operations, starts with `socket`. With $\bar{N} = 5$, PST was able to detect the malicious executions because of the unusual sequence. However, with $\bar{N} = 3$, it was *not* able to detect the extra FTP operation on Flow 1 at all. This is because the sequence `close-munmap-write-socket` is a legitimate one used at the boundary between File Write and the FTP upload stages (which is described in Figure 2 in Section 1). That is, the short sequence of calls was not able to differentiate the sequence from the ones generated by the transition between the legitimate FTP and the extra FTP stages. For Flow 2, PST was not able to detect the attacks at all for any \bar{N} (in addition to 3 and 5). This is because, as explained in Section 6.2, these executions are identical to the legitimate executions on Flow 1. Hence, both SCFD and PST were not able to catch the attacks.
3. **Attack 3 (Data corruption):** This attack only alters the number of `read` and `write` calls. In particular, the size of modified images affects the length of `write` chain in `mmap-write-...-write-close` sequence made when writing the compressed JPEG image to a file (see ‘File Write’ stage in Figure 2). When an attacker modifies the image, the length of the `write` chain becomes shorter. However, PST with a short sequence ($\bar{N} = 3$) was not able to classify any of the executions as being malicious because the length of the chain becomes 4 when the image is corrupted. Hence, the malicious executions were detected with $\bar{N} = 5$, because the probability of seeing `close` after `mmap-write-write-write-write` is zero given the training set. Note that, if the image size got larger instead and thus made the `write` chain longer than usual, sequence-based methods cannot detect this behavior because the only change would be that there are more chains of `write` that have a legitimate length. SCFD, on other hand, can detect this type of attack easily.
4. **Attack 4 (Shellcode execution):** Attack 4 was easily caught by both methods because `execve` call was never seen in the normal trace.

The results suggest that sequence-based approaches (the PST method in our evaluation) are sensitive to local, temporal variations, *e.g.*, an unusual transition from `write` to `write` instead of to `stat`. Our SCFD might not catch such a small, local variation.¹² However, sequence-based approaches fail to detect abnormal deviations in high-level, naturally variable execution contexts such as network activities or diverse data. This is because these require a global view on the frequencies of different system call types made during an entire execution. Hence, one can use these two approaches together to improve the overall accuracy of the system call-based IDS for embedded systems. From the implementation perspective, one may apply a sequence-based method to each system call observed and at the same

¹¹Some of the executions included FTP server error (due to too many connections) that caused the malicious FTP session to be disconnected. Both SCFD and PST classified such executions to be malicious and hence we excluded them from the accuracy results.

¹²In fact it can because having the extra `stat` is a strong indicator of anomaly by our normal execution model shown in Table 1.

Table 3: Time Complexity of our Analysis

# of system call types	Number of instructions	Avg. (Stdev.) of analysis times
5	2175	0.914 μs (0.553 μs)
10	4875	2.624 μs (1.405 μs)
14	8125	5.231 μs (1.965 μs)

time create an SCFD during an entire execution. Then, at the execution boundary, the SCFD method can be applied to check if the high-level execution context is anomalous or not.

6.4 Time Complexity

To evaluate the time complexity of the proposed detection method, we measured the number of instructions retired by the function that finds the closest cluster (Eq. (3)) among the five clusters given a new observation and the average time to perform the analysis.¹³

As Table 3 shows, the detection process is fast. This is possible because we store Σ^{-1} , the inverse of the covariance matrix, of each cluster, not Σ . A Mahalanobis distance is calculated in $\mathcal{O}(D^2)$, where D is the number of system call types being monitored, since in $(\mathbf{x}^* - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x}^* - \boldsymbol{\mu})$, the first multiplication takes $\mathcal{O}(D^2)$ and the second one takes $\mathcal{O}(D)$. Note that it would have taken $\mathcal{O}(D^3)$ if we stored the covariance matrix itself instead of its inverse; since a $D \times D$ matrix inversion takes $\mathcal{O}(D^3)$.

Note, again, that the monitoring and detection methods are *not in the critical path*, *i.e.*, they do not affect the execution of the applications we monitor since they are *offloaded* onto the secure core. More importantly, the time complexity of our method is *independent* of how often and many times the application uses system calls; it only depends on the number of system call types being monitored. This is determined in the training phase and does not change during the monitoring phase (see Section 3.4). On the other hand, the overheads of sequence-based approaches are highly dependent on the application complexity (*i.e.*, how many system calls are made). Hence, the deterministic time complexity of our SCFD method makes it particularly suitable for embedded systems.

6.5 Limitations and Possible Improvements

One of the limitations of our detection algorithm is that it checks for intrusions *after* execution is complete (at least for that invocation). Thus, if an attack tries to suddenly break the system, we cannot detect or prevent it. Combining a sequence-based method with our SCFD can be a solution if such attacks can be detectable by the former. If not, one can increase the chances of detection such problems by splitting the whole execution range into blocks [45] and checking for the distribution of system calls made in each block as soon as the execution passes each block boundary. This, however, would need more computation in the secure core at run-time, more storage in the SPM and a few more code modifications.

Another way to handle this problem is to combine this analysis/detection with other behavioral signals, especially ones that have a finer granularity of checks, *e.g.*, timing [45]. Since some blocks may use very few system calls (perhaps none) or even a very stable subset of such calls we can monitor the execution time spent in such a block to reduce the SCFD-based overheads (which is still low). This keeps the profile from bloating and prevents the system from having to carry out the legitimacy tests. We can also use the timing information in conjunction with the system call distribution; *i.e.*, by learning the normal time to execute a distribution of system calls, we can enforce a policy where each application block executes all of its system calls within (fairly) tight ranges. This is, of course,

¹³Simics is not a cycle-accurate simulator. Thus, the times are measured on a real machine with Intel Core I5 1.3GHz dual-core processor. The analysis code is compiled with `-O0` option. The statistic is based on 10,000 samples.

provided that the system calls themselves do not show unpredictable timing behavior. This makes it much harder for an attacker who imitates system calls [43] or one that replaces certain system calls with malicious functions [41].

Our model represents the histograms of system calls. In fact, each signature (histogram) would be represented by a mixture of multinomial distributions. That is, our model assumes generative models which select a mixture of multinomial parameters and then generate a histogram of system calls. Here, we assume that each multinomial distribution can be accurately represented by the multivariate Gaussian distributions (the multivariate Gaussian approximation of multinomial for large numbers of samples). In this regard, the assumption on our model is closely related to the *topic model* such as Latent Dirichlet Allocation [4, 3]. Here, we build a pragmatic lightweight module. One of the main drawbacks of the k -means clustering algorithm is that one may need to know or predefine the number of clusters. That is, system behaviors should be correctly represented by k numbers of multinomial (Gaussian) distributions of histogram. Some large-scale systems would have many heterogeneous modes (distributions). In this case, the appropriate solutions would be using non-parametric topic models such as Dirichlet process. However, we empirically observed that many embedded systems with predictable behavior can be represented by a tractable number of clusters. Thus, we use a simpler model with the k -means cluster.

7 Related Work

Forrest *et al.* [14] build a database of look-ahead pairs of system calls; for each system call type, what is the next i^{th} system call for $i = 1, 2$, up to N . Then, given a longer sequence of length $L > N$, the percentage of mismatches is used as the metric to determine abnormality. Hofmeyr *et al.* [16] extends the method by profiling unique sequences of *fixed* length N , called an N -gram, to reduce the database size. The legitimacy test for a given sequence of length N is carried out by calculating the smallest Hamming distance between it and the all sequences in the database. The N -gram model requires a prior assumption on suitable N because it affects the accuracy as well as the database size. Marceau [29] proposes a finite state machine (FSM) based prediction model to relax these requirements and Eskin *et al.* [12] further improves by employing a *wild-card* for compact sequence representation. Markovian techniques such as Hidden Markov model (HMM) [44] and variable-order Markov chain [40] have also been explored. Chandola *et al.* [6] provides an extensive survey on various anomaly detection techniques for discrete sequences. A similar approach to our work is [5], in which the system call counts of Android applications (traced by a software tool called *strace*) are used to find malicious apps. Using a crowdsourcing, the approach collects the system call counts of a particular application from multiple users and applies k -means (with Euclidean distance metric) to divide them into *two* clusters. A smaller cluster is considered to be malicious based on the assumption that benign apps are the majority.

There has also been work on system call arguments monitoring. Mutz *et al.* [32] introduce several techniques to test anomalies in argument lengths, character distribution, argument grammar, *etc.* Maggi *et al.* [26] use a clustering algorithm to group system call invocations that have similar arguments.

As previously mentioned, the usual way of system call instrumentation relies on an audit module in the OS layer. Hardware-based system call monitoring mechanism can improve the overall security of the system by cutting off a potential vulnerability – the software audit module. Pfoh *et al.*, [34] proposed Nitro, a hardware-based system call tracing system where system calls made inside virtual machines in manner similar to ours (Section 4.2). We note that our detection method (Section 3) is orthogonal to how system calls are traced. Hence we can implement it on systems like Nitro. Other types of instrumentation include static analysis of program source code [42] and user-level processes for system call interposition [17].

The SecureCore architecture [45] takes advantage of the *redundancy* of a multicore processor; a *secure core* is used to monitor the run-time execution behavior of target applications running on a *monitored*

core. The original architecture is designed to watch applications *timing* behavior. [46] extends the architecture by building a memory behavior monitoring framework for system-wide anomaly detection in real-time embedded systems. There also exists some work in which a multicore processor (or a coprocessor) is employed as a security measure, such as [7, 20, 9, 30, 25] for instruction-grain monitoring.

8 Conclusion

In this paper we presented a lightweight intrusion detection method that uses application execution contexts learned from system call frequency distributions of embedded applications. We demonstrated that the proposed detection mechanism could effectively complement sequence-based approaches by detecting anomalous behavior due to changes in high-level execution contexts. We also proposed certain architectural modifications to aid in the monitoring and analysis process. The approaches presented in the paper are limited in terms of the target applications and demonstration. Hence, as future work, we intend to implement the proposed architecture on a soft processor core [22] and to evaluate our method with real-world applications. We also plan to improve the learning and analysis methods using the topic modeling approach (explained in Section 6.5) to deal with large-scale heterogeneous behaviors of complex embedded applications.

References

- [1] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti. Control-flow integrity principles, implementations, and applications. *ACM Trans. Inf. Syst. Secur.*, 13(1):4:1–4:40, Nov. 2009.
- [2] R. Begleiter, R. El-Yaniv, and G. Yona. On prediction using variable order markov models. *J. Artif. Int. Res.*, 22(1):385–421, Dec. 2004.
- [3] D. M. Blei. Probabilistic topic models. *Commun. ACM*, 55(4):77–84, Apr. 2012.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, 2003.
- [5] I. Burguera, U. Zurutuza, and S. Nadjm-Tehrani. Crowdroid: Behavior-based malware detection system for android. In *Proc. of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, 2011.
- [6] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection for discrete sequences: A survey. *IEEE Trans. on Knowl. and Data Eng.*, 24(5):823–839, 2012.
- [7] S. Chen, M. Kozuch, T. Strigkos, B. Falsafi, P. B. Gibbons, T. C. Mowry, V. Ramachandran, O. Ruwase, M. Ryan, and E. Vlachos. Flexible hardware acceleration for instruction-grain program monitoring. In *Proc. of the International Symposium on Computer Architecture*, 2008.
- [8] J. Criswell, N. Dautenhahn, and V. Adve. Kcofi: Complete control-flow integrity for commodity operating system kernels. In *Proc. of the IEEE Symposium on Security and Privacy*, 2014.
- [9] D. Y. Deng, D. Lo, G. Malysa, S. Schneider, and G. E. Suh. Flexible and efficient instruction-grained run-time monitoring using on-chip reconfigurable fabric. In *Proc. of the IEEE/ACM International Symposium on Microarchitecture*, pages 137–148, 2010.
- [10] D. E. Denning. An intrusion-detection model. *IEEE Trans. Softw. Eng.*, 13(2):222–232, Feb. 1987.
- [11] e600 Core Reference Manual. http://www.freescale.com/files/32bit/doc/ref_manual/E600CORERM.pdf.

- [12] E. Eskin. Modeling system calls for intrusion detection with dynamic window sizes. In *Proc. of DARPA Information Survivability Conference and Exposition II*, 2001.
- [13] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(7):179–188, 1936.
- [14] S. Forrest, S. A. Hofmeyr, A. Somayaji, and T. A. Longstaff. A sense of self for unix processes. In *Proc. of the IEEE Symposium on Security and Privacy*, 1996.
- [15] M. Handley, V. Paxson, and C. Kreibich. Network intrusion detection: evasion, traffic normalization, and end-to-end protocol semantics. In *Proc. of the conference on USENIX Security Symposium*, 2001.
- [16] S. A. Hofmeyr, S. Forrest, and A. Somayaji. Intrusion detection using sequences of system calls. *J. Comput. Secur.*, 6(3):151–180, 1998.
- [17] K. Jain and R. Sekar. User-level infrastructure for system call interposition: A platform for intrusion detection and confinement. In *Proc. Network and Distributed Systems Security Symposium*, 1999.
- [18] I. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics, 2002.
- [19] K. Kanistrasy, G. Martinsy, M. J. Rutherfordz, and K. P. Valavanis. A survey of unmanned aerial vehicles (uavs) for traffic monitoring. In *Proc. of the International Conference on Unmanned Aircraft Systems*, 2013.
- [20] H. Kannan, M. Dalton, and C. Kozyrakis. Decoupling dynamic information flow tracking with a dedicated coprocessor. In *Proc. of the IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 105–114, 2009.
- [21] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental security analysis of a modern automobile. In *Proc. of the IEEE Symposium on Security and Privacy*, 2010.
- [22] LEON3 Processor. <http://www.gaisler.com/index.php/products/processors/leon3>.
- [23] A. Likas, N. Vlassis, and J. J. Verbeek. The global k-means clustering algorithm. *Pattern Recognition*, 36(2):451 – 461, 2003.
- [24] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [25] D. Lo, M. Ismail, T. Chen, and G. E. Suh. Slack-aware opportunistic monitoring for real-time systems. In *Proc. of the IEEE Real-Time and Embedded Technology and Applications Symposium*, 2014.
- [26] F. Maggi, M. Matteucci, and S. Zanero. Detecting intrusions through system call sequence and argument analysis. *IEEE Transactions on Dependable and Secure Computing*, 7(4):381–395, 2010.
- [27] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hillberg, J. Hgberg, F. Larsson, A. Moestedt, and B. Werner. Simics: A full system simulation platform. *IEEE Computer*, 35(2):50–58, 2002.
- [28] P. C. Mahalanobis. On the generalized distance in statistics. *Proc. of the National Institute of Sciences*, 2:49–55, 1936.
- [29] C. Marceau. Characterizing the behavior of a program using multiple-length n-grams. In *Proc. of the workshop on New security paradigms*, 2000.

- [30] S. Mohan, S. Bak, E. Betti, H. Yun, L. Sha, and M. Caccamo. S3A: Secure system simplex architecture for enhanced security and robustness of cyber-physical systems. In *Proc. of the ACM International Conference on High Confidence Networked Systems*, 2013.
- [31] Freescale MPC8641D Integrated Host Processor. http://www.freescale.com/files/32bit/doc/ref_manual/MPC8641DRM.pdf.
- [32] D. Mutz, F. Valeur, G. Vigna, and C. Kruegel. Anomalous system call detection. *ACM Trans. Inf. Syst. Secur.*, 9(1):61–93, Feb. 2006.
- [33] G. Parmer and R. West. Hijack: Taking control of cots systems for real-time user-level services. In *Proc. of the IEEE Real-Time and Embedded Technology and Applications Symposium*, 2007.
- [34] J. Pfoh, C. Schneider, and C. Eckert. Nitro: hardware-based system call tracing for virtual machines. In *Proc. of the International conference on Advances in information and computer security*, 2011.
- [35] Linux/x86-64 shellcode - add user with passwd. <http://www.shell-storm.org/shellcode/files/shellcode-801.php>.
- [36] Linux/PPC shellcode - execve() of /bin/sh. <http://www.shell-storm.org/shellcode/files/shellcode-89.php>.
- [37] D. Shepard, J. Bhatti, and T. Humphreys. Drone hack: Spoofing attack demonstration on a civilian unmanned aerial vehicle. *GPS World*, Aug 2012.
- [38] SNARE: System iNtrusion Analysis and Reporting Environments. <http://http://www.intersectalliance.com/>.
- [39] R. Sommer and V. Paxson. Outside the closed world: On using machine learning for network intrusion detection. In *Proc. of the IEEE Symposium on Security and Privacy*, 2010.
- [40] P. Sun, S. Chawla, and B. Arunasalam. Mining for outliers in sequential databases. In *Proc. of the SIAM International Conference on Data Mining*, 2006.
- [41] Hijacking system calls with loadable kernel modules. <http://r00tkit.me/?p=46>.
- [42] D. Wagner and D. Dean. Intrusion detection via static analysis. In *Proc. of the IEEE Symposium on Security and Privacy*, 2001.
- [43] D. Wagner and P. Soto. Mimicry attacks on host-based intrusion detection systems. In *Proc. of the ACM conference on Computer and communications security*, 2002.
- [44] C. Warrender, S. Forrest, and B. Pearlmutter. Detecting intrusion using system calls: alternative data models. In *Proc. of the IEEE Symposium on Security and Privacy*, 1999.
- [45] M.-K. Yoon, S. Mohan, J. Choi, J.-E. Kim, and L. Sha. SecureCore: A multicore-based intrusion detection architecture for real-time embedded systems. In *IEEE Real-Time Embedded Technology and Applications Symposium*, 2013.
- [46] M.-K. Yoon, S. Mohan, J. Choi, and L. Sha. Memory Heat Map: Anomaly detection in real-time embedded systems using memory behavior. In *Proc. of the ACM/EDAC/IEEE Design Automation Conference*, 2015.
- [47] M. M. Z. Zadeh, M. Salem, N. Kumar, G. Cutulenco, and S. Fischmeister. SiPTA: Signal processing for trace-based anomaly detection. In *Proc. of the International Conference on Embedded Software*, 2014.

APPENDIX

A Cutoff Distance

In general, there is no analytic solution for calculating the cumulative distribution function (CDF) for multivariate normal distributions. However, it is possible to derive the CDF with Mahalanobis distance. The cutoff distance θ can be derived by finding the smallest distance that makes the probability that a data point \mathbf{x} , which in fact belongs to the cluster and has a distance farther than θ , is not greater than $p_0 = 0.01$ or 0.05 .

First, let z be a Mahalanobis distance from a multivariate normal distribution. Then,

$$\int_0^\theta c \cdot e^{-\frac{1}{2}z^2} dz = 1 - p_0, \quad (5)$$

where c is a normalizing constant that satisfies Eq. (5) with $\theta = \infty$ and $p_0 = 0$ by the definition of a probability density function. This results in $c = 1/1.25331$ because

$$\int_0^\infty e^{-\frac{1}{2}z^2} dz \simeq \left[1.25331 \cdot \text{erf}(0.707107 \cdot z) \right]_0^\infty = 1.25331,$$

where $\text{erf}(z)$ is the *error function* and is 1 and 0 for $z = \infty$ and $z = 0$, respectively. Accordingly, Eq. (5) becomes

$$\begin{aligned} \frac{1}{1.25331} \int_0^\theta e^{-\frac{1}{2}z^2} dz &\simeq \frac{1}{1.25331} \left[1.25331 \cdot \text{erf}(0.707107 \cdot z) \right]_0^\theta \\ &= \text{erf}(0.707107 \cdot \theta) = 1 - p_0. \end{aligned}$$

Therefore, the cutoff distance θ for a significant level p_0 is

$$\theta = \frac{\text{erf}^{-1}(1 - p_0)}{0.707107}. \quad (6)$$

For $p_0 = 1\%$ and 5% , $\theta \sim 2.57583$ and 1.95996 , respectively. Figure 9 shows the cutoff distance for $0\% \leq p_0 \leq 100\%$. The cutoff distance is not bounded (*i.e.*, $\theta = \infty$) when $p_0 = 0\%$ and is 0 when $p_0 = 100\%$.




Figure 9: The cutoff distance θ for significant level p_0 .